Defense Technical Information Center
Compilation Part Notice

ADP014341

TITLE: Quantum Computing: From Bragg Reflections to Decoherence Estimates

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: Materials Research Society Symposium Proceedings. Volume 746. Magnetoelectronics and Magnetic Materials - Novel Phenomena and Advanced Characterization

To order the complete compilation report, use: ADA418228

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, etc. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:
ADP014306 thru ADP014341

Q8.5

QUANTUM COMPUTING: FROM BRAGG REFLECTIONS TO DECOHERENCE ESTIMATES

Peter Pfeifer and Chen Hou
Department of Physics, University of Missouri
Columbia, MO 65211, U.S.A.

ABSTRACT

We give an exposition of the principles of quantum computing (logic gates, exponential parallelism from polynomial hardware, fast quantum algorithms, quantum error correction, hardware requirements, and experimental milestones). A compact description of the quantum Fourier transform to find the period of a function—the key step in Shor's factoring algorithm—illustrates how parallel state evolution along many classical computational paths produces fast algorithms by constructive interference similar to Bragg reflections in x-ray crystallography. On the hardware side, we present a new method to estimate critical time scales for the operation of a quantum computer. We derive a universal *upper* bound on the probability of a computation to fail due to decoherence (entanglement of the computer with the environment), as a function of time. The bound is parameter-free, requiring only the interaction between the computer and the environment, and the time-evolving state in the absence of any interaction. For a simple model we find that the bound performs well and decoherence is small when the energy of the computer state is large compared to the interaction energy. This supports a recent estimate of minimum energy requirements for quantum computation.

INTRODUCTION

A quantum computer puts to use the fact that a quantum particle, such as an electron passing through a double-slit apparatus or an atom or molecule traversing a matter-wave interferometer [1], can be in two different locations at the same time. By equating different locations—as a paradigm, we take an electron in the lowest orbit or in an excited orbit of an atom—to binary digits 0 and 1, one may interpret the time-evolving state of the particles as executing several computations at the same time. One set of locations at a given time describes the result of one computation. Thus one atom can do two computations at once; two atoms can do four; three atoms can do eight. The challenge is to coerce the atoms to follow trajectories that amount to meaningful computations and to read out a definite result from the multitude of computations occurring in parallel. The control of trajectories is the hardware part of the challenge; the design of useful trajectories—algorithms that are superior to classical algorithms—is the software part of the challenge.

This paper is a comprehensive survey of both aspects. It illustrates in terms of explicit case studies the key components of a quantum computer, namely, the design of a computational trajectory (Shor's algorithm), the trajectory's superior performance (exponential parallelism), and the reliable operation of the computer vis-à-vis decoherence, with emphasis on analogies to familiar physical phenomena. The case studies apply regardless of whether the computer is implemented in terms of nuclear spins, ion traps, cavity quantum electrodynamics, superconducting currents, or spintronics. General references are [2-9].

QUANTUM VERSUS CLASSICAL COMPUTATION

A classical computer manipulates strings of N classical bits, $(n_1, ..., n_N)$ with $n_j = 0$ or 1 $(j = 1, ..., N)$, in such a way that intermediate states of the computation are also strings of classical bits. A quantum computer manipulates states of N two-level atoms or any other two-level entities, $|n_1, ..., n_N\rangle$ with $n_j = 0$ if the j-th atom is in the ground state and $n_j = 1$ if it is in the excited state, in such a way that intermediate states are superpositions of the states $|n_1, ..., n_N\rangle$. The 2^N states $|n_1, ..., n_N\rangle$ ("computational basis") are product states in which each atom is either in the ground state or excited state, and n_j is called the value of the j-th qubit ("quantum bit"); they represent the strings of classical bits. The superpositions include states in which an atom no longer has a sharp value of n_j (indefinite bit value), and states in which an atom no longer exists in a state separate from the other atoms (entangled state); both have no classical counterpart.

A quantum computation starts with a product state $|n_1, ..., n_N\rangle$, lets the state evolve according to the Schrödinger equation, $i\hbar \frac{d}{dt}|\psi(t)\rangle = H(t)|\psi(t)\rangle$, with initial condition $|\psi(0)\rangle = |n_1, ..., n_N\rangle$ and time-dependent Hamiltonian $H(t)$ driving the coupled atoms, and ends with the measurement of the values of the qubits of the state $|\psi(t)\rangle$. The Hamiltonian generates the unitary time-evolution operator $U(t)$ which takes the initial state into the final state,

$$|\psi(t)\rangle = T\exp\left(-(i/\hbar)\int_0^t H(s)ds\right)|\psi(0)\rangle = U(t)|\psi(0)\rangle \qquad (1)$$

with T the time-ordering operator. The measurement transforms $|\psi(t)\rangle$ into the output state $|n_1', ..., n_N'\rangle$ with probability $|\langle n_1', ..., n_N'|\psi(t)\rangle|^2$. The output is probabilistic because quantum measurements are so. The computation $|n_1, ..., n_N\rangle \mapsto |\psi(t)\rangle$ is a unitary transformation, hence reversible; the readout $|\psi(t)\rangle \mapsto |n_1', ..., n_N'\rangle$ is a projection ("collapse of the wave function"), hence irreversible. Thus, to perform a specific computation, one must drive the atoms with a specific Hamiltonian; to read out the result, one must send the atoms through a series of state detectors.

A quantum computer is more powerful than a classical computer for two reasons. (i) The quantum state space is much larger than the classical state space: N qubits can be in an infinite number of different states (any point on the unit sphere of the complex Hilbert space spanned by the 2^N basis vectors $|n_1, ..., n_N\rangle$); N classical bits can only be in 2^N different states (the points where the 2^N coordinate axes intersect the unit sphere, Fig. 1). If the expansion coefficients of the superpositions can be controlled with accuracy ε ($\varepsilon \ll 1$), the sphere hosts $O(\varepsilon^{-(2^N-1)})$ distinct states. Thus a quantum computer can store and access an exponentially large number of states compared to a classical computer. (ii) The quantum computer operates in a massively parallel way: if the initial state is the uniform superposition of all basis states,

$$|\psi(0)\rangle = 2^{-N/2} \sum_{n_1, ..., n_N = 0,1} |n_1, ..., n_N\rangle, \qquad (2)$$

the time evolution computes simultaneously $U(t)|n_1, ..., n_N\rangle$ for all 2^N possible inputs $|n_1, ..., n_N\rangle$ by linearity of $U(t)$. The matrix element $\langle n_1', ..., n_N'|U(t)|n_1, ..., n_N\rangle$ is the probability amplitude that the computation converts the input $|n_1, ..., n_N\rangle$ into the output $|n_1', ..., n_N'\rangle$, along all possible classical computational paths in parallel (Feynman's path integral). A classical computation can follow only a single path. The aim is to choose $U(t)$ so that computational paths of no interest

Figure 1. A classical bit can be in state $|0\rangle$ or $|1\rangle$ (left). A quantum bit can be in any superposition $a|0\rangle + b|1\rangle$, where the complex numbers a and b satisfy $|a|^2 + |b|^2 = 1$ (right). The states are shown as expectation values of the spin vector.

cancel each other by destructive interference, and paths of interest add constructively (selection of relevant computations by quantum interference, Fig. 2).

Any computation $U(t)$, also called a quantum circuit, can be approximated by sequential application of a finite set of unitary transformations that operate on only one or two qubits. An example for such a "universal set of quantum logic gates" is

$$H_j = \left(|0_j\rangle\langle 0_j| + |0_j\rangle\langle 1_j| + |1_j\rangle\langle 0_j| - |1_j\rangle\langle 1_j|\right)/\sqrt{2} \quad \text{("Hadamard gate")}, \tag{3}$$

$$T_j = |0_j\rangle\langle 0_j| + e^{i\pi/4}|1_j\rangle\langle 1_j| \quad \text{("T gate")}, \tag{4}$$

$$C_{jk} = |0_j0_k\rangle\langle 0_j0_k| + |0_j1_k\rangle\langle 0_j1_k| + |1_j0_k\rangle\langle 1_j1_k| + |1_j1_k\rangle\langle 1_j0_k| \quad \text{("controlled-not gate")}, \tag{5}$$

where $|n_j\rangle\langle n'_j|$ acts only on qubit j, and $|n_jn_k\rangle\langle n'_jn'_k|$ only on qubits j and k. They correspond to the logic gates in a classical computer, but are reversible (the classical 'and' and 'exclusive or' gates are irreversible). The Hadamard gate transforms the states $|0_j\rangle$ and $|1_j\rangle$ into the superpositions $(|0_j\rangle \pm |1_j\rangle)/\sqrt{2}$; the T gate shifts the phase of the excited state relative to the ground state by $\pi/4$; and the controlled-not gate flips the "target" qubit k if and only if the "control" qubit j is in the excited state. The three gates are the analog of an optical beam splitter, phase shifter (refractive medium), and conditional mirror, respectively. Only the Hadamard gate creates multiple computational paths; the other two transform a single basis state into a single basis state. To approximate a general $U(t)$ to within accuracy ε requires $O(4^N N^2 [\ln(4^N N^2/\varepsilon)]^\alpha)$ gates, where $\alpha \approx 2$ (Solovay-Kitaev theorem [6]) and the leading factor 4^N is set by the number of matrix elements of $U(t)$. For special computations, however, often a much smaller number of gates, independent of ε, suffices, as we shall see in the next section.

Figure 2. Application of the Hadamard gate to state $|n_j\rangle$ creates two computational paths, one leading to $|0_j\rangle$ with probability amplitude $1/\sqrt{2}$, the other to $|1_j\rangle$ with amplitude $(-1)^n/\sqrt{2}$. Application of the Hadamard gate twice to $|0_j\rangle$ creates four paths, leading to $|0_j\rangle$ with amplitude $[(+1)(+1)+(+1)(+1)]/2$, and to $|1_j\rangle$ with amplitude $[(+1)(+1)+(+1)(-1)]/2$ (cancellation of trajectories leading to $|1_j\rangle$).

AN EXPLICIT EXAMPLE—FINDING THE PERIOD OF A FUNCTION

A widely used class of encryption systems for public transmission of sensitive data (RSA public key cryptosystem [10]) derive their security from the difficulty of factoring a large, publicly transmitted integer. The fastest known classical algorithm factors an N-bit number in time $O(2^{const \times N^{1/3}(\ln N)^{2/3}})$. Shor's celebrated quantum algorithm [11, 12] for the same task requires only time $O(N^2(\ln N)(\ln \ln N))$. The key task in Shor's algorithm is to find the period r of a periodic function $f(x)$ related to the number to be factored, and the tool used to determine the period is the discrete Fourier transform. The values of the Fourier variable at which the transform does not vanish yield the period, just as in x-ray diffraction from a crystal the scattered wave vectors at which the x-ray intensity is nonzero yield the crystal periodicity (Bragg reflection at wave vectors equal to a reciprocal lattice vector).

The discrete Fourier transform assigns to a given complex-valued function $f(x)$ (not necessarily periodic) the function

$$\hat{f}(y) = 2^{-N/2} \sum_{x=0}^{2^N-1} e^{2\pi i x y/2^N} f(x), \tag{6}$$

where both x and y run through the integers $0, 1, ..., 2^N - 1$. Its quantum computation reads

$$\sum_{y=0}^{2^N-1} \hat{f}(y)|y_1, ..., y_N\rangle = U(t) \sum_{x=0}^{2^N-1} f(x)|x_1, ..., x_N\rangle, \tag{7}$$

$$U(t)|x_1, ..., x_N\rangle = 2^{-N/2} \sum_{y=0}^{2^N-1} e^{2\pi i x y/2^N} |y_1, ..., y_N\rangle, \tag{8}$$

$$x = x_1 2^{N-1} + x_2 2^{N-2} + ... + x_N 2^0, \tag{9a}$$

$$y = y_1 2^{N-1} + y_2 2^{N-2} + ... + y_N 2^0. \tag{9b}$$

Equation (8) defines the operator $U(t)$ in terms of the computational basis; (9) is the binary expansion of the summation variables in (7, 8); and Eq. (7) is established by substitution of (6) and (8). The sums in (7), running over the whole basis, completely describe the function and its Fourier transform. The implementation of $U(t)$ in terms of logic gates is given by

$$U(t)|x_1, ..., x_N\rangle = 2^{-N/2} \sum_{y=0}^{2^N-1} e^{2\pi i xy/2^N} |y_1, ..., y_N\rangle$$

$$= 2^{-N/2} \prod_{j=1}^{N} \left(|0_j\rangle + e^{2\pi i 0.x_{N-j+1}x_{N-j+2}...x_N} |1_j\rangle \right)$$

$$= R(H_1 B_{1,2} B_{1,3} \cdots B_{1,N})(H_2 B_{2,3} B_{2,4} \cdots B_{2,N}) \cdots (H_{N-1} B_{N-1,N})(H_N)|x_1, ..., x_N\rangle. \tag{10}$$

The first line recalls the action of $U(t)$ on a basis state; the second line shows that the result is a product state, not an entangled state (the binary fractions $0.x_{N-j+1}x_{N-j+2}...x_N$ arise from the expansion of $xy/2^N$); and the third line expresses the product in terms of successive applications of the Hadamard gate, the controlled-phase-shift gate,

$$B_{jk} = |0_j 0_k\rangle\langle 0_j 0_k| + |0_j 1_k\rangle\langle 0_j 1_k| + |1_j 0_k\rangle\langle 1_j 0_k| + e^{\pi i/2^{k-j}} |1_j 1_k\rangle\langle 1_j 1_k|, \tag{11}$$

and the qubit reversal operator, $R|n_1, ..., n_N\rangle = |n_N, ..., n_1\rangle$ (implemented by reading the qubits in reverse order), on the original basis state.

Equation (10) is the quantum fast Fourier transform. It is an explicit example for the decomposition of a quantum circuit into a product of one- and two-qubit operations and shows that the transform, without any approximation, can be carried out with only $N(N+1)/2$ operations. This is much less than the Solovay-Kitaev bound and much less than the $O(2^N N)$ operations in the classical fast Fourier transform. The N Hadamard gates create 2^N classical paths along which the state $|x_1, ..., x_N\rangle$ evolves in parallel.

The period r is obtained by measuring the qubits of state (7). The measurement yields the output state $|y_1, ..., y_N\rangle$ with probability

$$|\langle y_1, ..., y_N |\psi(t)\rangle|^2 = |\hat{f}(y)|^2. \tag{12a}$$

If r divides 2^N, the probability is

$$|\hat{f}(y)|^2 = \frac{2^N}{r^2} \left| \sum_{x=0}^{r-1} e^{2\pi i xy/2^N} f(x) \right|^2 \neq 0 \quad \text{if } y/2^N = 0, 1/r, 2/r, ..., (r-1)/r, \tag{12b}$$

$$|\hat{f}(y)|^2 = 0 \quad\quad\quad\quad \text{else,} \tag{12c}$$

from (6). Thus as advertised, only "good" y's, related to r, remain. Outcomes unrelated to r occur with probability zero, by virtue of destructive interference. The good y's yield r as the largest denominator of the reduced fractions $y/2^N$, as y runs through all outcomes $|y_1, ..., y_N\rangle$. The largest denominator from m distinct outcomes ($m \ll r$) equals r with probability better than $1 - (1 - 0.561/\ln\ln r)^m$ (prime number theorem), so even a period as large as 2^{100} is found with a success rate better than 99% from only 33 distinct outcomes. If r does not divide 2^N, destructive interference is not perfect, and the measurements with high probability yield y's satisfying $y/2^N = 0, 1/r, 2/r, ..., (r-1)/r$ approximately. In this case, r is found from a slightly more involved analysis, but still with only $O(\ln N)$ Fourier transforms and measurements. This gives the period in time $O(N^2 \ln N)$. With this efficiency, it turns out that the bottleneck in Shor's algorithm is the computation of $f(x)$, i.e., the preparation of the initial state $\sum_x f(x)|x_1, ..., x_N\rangle$, which we have suppressed here. The determination of the period, including the preparation of the initial state, requires time $O(N^2(\ln N)(\ln\ln N))$.

Table I. Exponential parallelism of the quantum Fourier transform from polynomial hardware. The second line is often expressed as that the quantum computer can do 2^N classical computations simultaneously. The abbreviations in the last column stand for 'effort' and 'payoff'.

	Number	Eq.	Effort/payoff
Logic gates acting on $|\psi(0)\rangle$	$N(N+1)/2$	(10)	Polynomial e.
Classical bit strings processed at once in $|\psi(t)\rangle$	2^N	(7)	Exponential p.
Classical paths followed at once in $|\psi(t)\rangle$	2^{2N}	(10)	Exponential p.
Distinct constructive interference events from $|\psi(t)\rangle$	$2^N r$	(12)	Exponential p.
Repeated computations of $|\psi(t)\rangle$ needed to find r	$O(\ln N)$	–	Logarithmic e.

Table I summarizes the algorithm in terms of parallel computing and interference of computational paths. The number of distinct constructive interference events increases with increasing period r, just as in x-ray diffraction the number of Bragg reflections increases with increasing size of the unit cell. The degree of entanglement of computational states varies in the opposite direction: If f is zero everywhere except at position x, then f has period $r = 2^N$, the initial state is a product state by $|\psi(0)\rangle = |x_1, ..., x_N\rangle$, and so is the final state $|\psi(t)\rangle$ by (10). In this case, neither the initial nor final state is entangled. If $f(x) = 2^{-N/2}$ for all x, then f has period $r = 1$, the initial state equals the state (2), and the final state is $|\psi(t)\rangle = |0, ..., 0\rangle$. In this case, the initial state is maximally entangled, and the final state is not entangled. This suggests that the smaller the period, the more entangled are relevant states in the computation.

OTHER FAST QUANTUM ALGORITHMS

Factoring integers can be reduced to finding integer solutions of Pell's equation, $x^2 - ay^2 = 1$, where a is a positive non-square integer; but a converse is not known. So it is of interest that a quantum algorithm exists also for Pell's equation. The algorithm is due to Hallgren [13] and finds the solution in polynomial time as in Shor's algorithm, instead of exponential time as on a classical computer.

Other quantum algorithms that outperform classical algorithms by orders of magnitude are: Grover's algorithm for "finding a needle in a haystack" (search of an item in a database of 2^N items in $O(2^{N/2})$ instead of $O(2^N)$ steps); estimation of the median and mean of 2^N items to precision ε in $O(N/\varepsilon)$ instead of $O(N/\varepsilon^2)$ steps; search of the minimum of a function sampled at 2^N points in $O(2^{N/2})$ instead of $O(2^N)$ steps; search of two distinct pre-images giving the same image of a two-to-one function sampled at 2^N points, in $O(2^{N/3})$ instead of $O(2^{N/2})$ steps; the Deutsch-Jozsa algorithm to determine if 2^N numbers are either all 0 ("constant function"), or half are 0 and half are 1 ("balanced function"), in one instead of up to $2^{N-1} + 1$ steps; and various allocation tasks and game-theoretic strategies.

Quantum algorithms for solving complex physical problems in polynomial time instead of exponential time include [9]: path integration with respect to Wiener measure; quantum random walks; simulation of the quantum baker's map; quantum lattice-gas model for the time-dependent Schrödinger equation of many-body systems; and the dynamical sign problem in fermionic systems.

QUANTUM ERROR CORRECTION

Noise from imperfect computer operation poses no fundamental barrier to large-scale computations. A quantum error-correction code encodes the N "logical qubits" into N' "carrier qubits" ($N' > N$), runs the carrier qubits through a group of accordingly encoded logic gates, transforms the noisy state by appropriate projection operators ("syndrome measurements") and unitary operators ("recovery of original carrier qubits") into an error-corrected state, and feeds the state into the next group of gates. Such periodic error correction prevents accumulation of errors in the state ("quantum Zeno effect"). At the end of the computation, the carrier qubits are decoded. The encoding spreads the state of the N logical qubits over all N' carrier qubits so that when the syndromes are measured, no information about the state of the logical qubits is revealed ("noise-less" or "decoherence-free" subspaces): the projections preserve superpositions of the logical qubits, and the state of the carrier qubits is highly entangled even if the logical-qubit state is not. Remarkably, a discrete set of corrections can correct a continuum of errors. For a code to correct any error on any M carrier qubits, a necessary condition is

$$N' \geq 4M + N \tag{13}$$

(Knill-Laflamme bound), and a sufficient condition for large N is

$$N/N' < 1 - 2\left[-x\log_2 x - (1-x)\log_2(1-x)\right]_{x=2M/N'} \tag{14}$$

(Gilbert-Varshamov bound). E.g., a code exists which encodes one logical qubit ($N = 1$) into 5 carrier qubits ($N' = 5$) and corrects any error on any one carrier qubit ($M = 1$) with 16 pairs of syndrome measurements and recovery operations (equality in the Knill-Laflamme bound).

Different codes require different encoding of gates, and of interest are encodings for which an error in an input carrier qubit or the gate operation propagates only to a small number of output carrier qubits. Specifically, an encoded gate is called fault-tolerant if a failure with probability p of any single component (e.g., one of the 10 "wires" feeding two logical qubits, each encoded by 5 carrier qubits, into a controlled-not gate) introduces an error in two or more carrier qubits in any logical output qubit with probability cp^2 at most, with c a constant and p small. Such a gate, when followed by error correction with $M = 1$, yields an error-free output with probability $1 - cp^2$, i.e., reduces the error probability from p to cp^2 if $p < 1/c$. Fault-tolerant Hadamard, T, and controlled-not gates exist with $c \approx 10^4$ and $d \approx 10^2$, where d is the number of operations on carrier qubits needed to encode and error-correct the gate. By hierarchical fault-tolerant encoding of all gates, a computation involving L gates can be carried out to within accuracy ε using only $O(L[\ln(L/\varepsilon)]^{\log_2 d})$ operations on carrier qubits, if $p < 1/c$ (threshold theorem for quantum computation). Thus, if the noise in individual carrier qubits is low enough, $p < 1/c \approx 10^{-4}$, arbitrarily large computations can be performed because the overhead for error correction grows only polynomial-logarithmically with the size of the computation, L.

EXPERIMENTAL STATE OF THE ART AND HARDWARE REQUIREMENTS

A remarkable array of experimental realizations of quantum computing devices and algorithms have been achieved to date. Current record holders with respect to the number of qubits

that can be controlled and prepared in well-defined states include, in the different categories of experimental implementation:

(i) 7-qubit nuclear magnetic resonance devices (NMR): three ^{1}H and four ^{13}C nuclei in *trans*-crotonic acid, each in the spin up or down state, with a total of twelve 2-qubit gates driven by radio-frequency pulses, preparing the "Schrödinger-cat state" $(|0000000\rangle + |1111111\rangle)/\sqrt{2}$ and converting it into the state $|0000000\rangle$ [14]; and five ^{19}F and two ^{13}C nuclei in a custom-made molecule, fully implementing, with a sequence of 300 radio-frequency pulses, Shor's algorithm to factor the number 15 [15];

(ii) a 4-qubit ion-trap device (IT): four ^{9}Be$^+$ ions in a linear electromagnetic trap, each in one of two hyperfine Zeeman levels and driven by Raman transitions [16];

(iii) a 3-qubit device based on cavity quantum electrodynamics (CQED): three Rb atoms, each in one of two Rydberg states (principal quantum numbers 49, 50, and 51) and coupled to a cavity mode with zero or one photon, driven by microwave pulses [17]; and

(iv) 1-qubit devices based on macroscopic persistent-current states in superconductors [18].

Computations carried out, in addition to Shor's algorithm, include: Grover's algorithm on a 2-qubit NMR device (1998/99); dynamics of quantum harmonic and anharmonic oscillators on a 2-qubit NMR device (1999); the Deutsch-Jozsa algorithm on a 5-qubit NMR device (2000); finding the order of a permutation on a 5-qubit NMR device (2000); correction of any one-qubit error on a 5-qubit NMR device (2001); noiseless encoding of one logical qubit in a 3-qubit NMR device (2001); quantum lattice-gas treatment of the diffusion equation and Burgers equation on a 2-qubit NMR device (2002); quantum baker's map on a 3-qubit NMR device (2002); protection of an IT qubit from decoherence (two ^{9}Be$^+$ ions encoding one qubit (2001)); controlled-not gate on a 2-qubit IT device (one ^{9}Be$^+$ ion implementing a motional and an internal-state qubit (2002)); and Grover's algorithm in a 5-"qubit" *classical* optical cavity, demonstrating that the algorithm requires no entanglement (2002).

Any experimental realization faces three challenges:

(i) The device must be able to control the state of each qubit separately, while allowing neighboring qubits to interact with each other for the operation of two-qubit gates. Control of individual qubits is achieved by different chemical shifts in NMR, laser beams driving ions spatially separated by tens of μm in IT, and pulses addressing successive atoms traveling at spatial separation of several cm in CQED. Qubits interact via spin-spin coupling in NMR, a shared phonon state of the ions (excitation of the center-of-mass mode) in IT, and electric dipole coupling between each atom and the cavity mode in CQED.

(ii) The system must be switchable, so that interactions executing a prescribed sequence of gate operations (Hamiltonian $H(t)$) can be turned on and off by external control. Switching is done by magnetic field pulses in NMR, and laser pulses in IT and CQED.

(iii) The computer must be well isolated from the environment so that the decoherence time t_d, the time at which the computer and the environment depart significantly from a product state (entanglement of the computer with the environment), is long compared to t_g, the time it takes to operate a single gate. At time t_d, a generic qubit state $a|0\rangle + b|1\rangle$ will have degraded into the mixture $|a|^2|0\rangle\langle0| + |b|^2|1\rangle\langle1|$ (density matrix), which no longer contains the interference terms necessary for quantum computation. Good isolation is provided by long spin-spin and spin-lattice relaxation times in NMR ($t_d = 10^{-2} - 10^8$ s, $t_g = 10^{-6} - 10^{-3}$ s), long-lived hyperfine levels and stable trap and laser operation in IT ($t_d = 10^{-1} - 10^0$ s, $t_g = 10^{-7} - 10^{-5}$ s), and the low spontaneous emission rate of Rydberg states and low photon escape rate from the cavity in CQED ($t_d = 10^{-3} - 10^0$ s, $t_g = 10^{-5} - 10^{-4}$ s).

DECOHERENCE ESTIMATES

The condition $t_g < t_d$ is a challenge because the computer must be a weakly open system for time t_g, but a strongly closed system for time t_d. The condition $p < 1/c$ in the threshold theorem for quantum computation is even more demanding. To see this, let $|\psi_i(t)\rangle$ be the state of the computer and environment interacting with each other, starting from a product state at $t = 0$; and let $|\psi_n(t)\rangle$ be the same state, but noninteracting (product state at all times):

$$|\psi_i(t)\rangle = T\exp\left(-(i/\hbar)\int_0^t \{H_c(s) + H_e(s) + V(s)\}ds\right)|\psi_c(0)\rangle|\psi_e(0)\rangle, \tag{15}$$

$$|\psi_n(t)\rangle = \left[T\exp\left(-(i/\hbar)\int_0^t H_c(s)ds\right)|\psi_c(0)\rangle\right]\left[T\exp\left(-(i/\hbar)\int_0^t H_e(s)ds\right)|\psi_e(0)\rangle\right]. \tag{16}$$

Here $H_c(t)$, $H_e(t)$, and $V(t)$ is the Hamiltonian of the computer, environment, and interaction between the two, respectively; and $|\psi_c(0)\rangle$ and $|\psi_e(0)\rangle$ is the initial state of the computer and environment, respectively. Suppose now the overlap ("fidelity"),

$$F(t) := \left|\langle\psi_i(t)|\psi_n(t)\rangle\right|, \tag{17}$$

decays exponentially with time. The decay constant is the inverse of the decoherence time, $F(t) = e^{-t/t_d}$, and the threshold condition requires the failure probability p to satisfy

$$p = 1 - F^2(t_g) = 1 - e^{-2t_g/t_d} < 1/c, \tag{18}$$

which gives the condition $t_g < t_d/(2c)$ for large c. Thus the computer must remain isolated from the environment for about 10^4 gate operations for sustained computations. Present devices can execute about 300 operations [15] and are far from this goal.

One may seek to reduce t_g and use (18) to estimate how short t_g should be for the threshold theorem to kick in. But for short times, the fidelity does not decay exponentially, and Eq. (18) is not appropriate. Also, the fidelity evolves differently if a gate is switched rapidly instead of slowly, which further limits the validity of (18). It is therefore of interest that a universal bound for the fidelity exists which suffers from none of these limitations. The bound reads

$$F(t) \geq F_-(t) := \cos\left(\min\left\{\hbar^{-1}\int_0^t \Delta(s)ds, \ \pi/2\right\}\right), \tag{19}$$

$$\Delta(s) := \sqrt{\langle\psi_n(s)|V^2(s)|\psi_n(s)\rangle - \langle\psi_n(s)|V(s)|\psi_n(s)\rangle^2}, \tag{20}$$

valid for all t. It requires only the knowledge of the product state (16) and interaction $V(s)$. The energy uncertainty $\Delta(s)$ is easy to evaluate compared to the task of computing the fidelity from (15-17). The bound is a "worst-case" decoherence estimate: if $F_-(t)$ is close to 1, so is $F(t)$ by $F(t) \leq 1$, and decoherence is *guaranteed* small. $F_-(t)$ equals 1 at $t = 0$, decreases with increasing t (the more slowly, the weaker the interaction), and reaches zero (trivial bound) when the integral exceeds $\pi\hbar/2$. The time τ at which $F_-(t)$ has dropped to the value $1/e$, $\hbar^{-1}\int_0^\tau \Delta(s)ds = 1.19$, gives a lower bound for the decoherence time, $\tau \leq t_d$ (Fig. 3). A particularly appealing form of inequality (19) is the *upper* bound

$$p(t) \le \sin^2 \left(\min \left\{ \hbar^{-1} \int_0^t \Delta(s)\, ds, \ \pi/2 \right\} \right)$$
(21)

for the probability of the quantum computation $|\psi_c(t)\rangle$ (first factor in (16)) to "fail" after time t due to imperfect isolation, $p(t) := 1 - F^2(t)$. The term "fail" is short for the interacting state to give possibly different results than the noninteracting state. The bound (21) estimates the failure probability rigorously, at any point in time, and without reference to gate-switching and decoherence times (parameter-free decoherence estimate), unlike (18).

The result (19) is an application of the variational principle of Pfeifer and Fröhlich [19] (see also [20]) which states that if $|\psi(t)\rangle$ is a solution of the time-dependent Schrödinger equation with Hamiltonian $H(t)$ and $|\psi_t(t)\rangle$ is a trial state with $|\psi_t(0)\rangle = |\psi(0)\rangle$, then

$$\left| \langle \psi(t) | \psi_t(t) \rangle \right| \ge \cos \left(\min \left\{ \hbar^{-1} \int_0^t \left\| \left(1 - |\psi_t(s)\rangle\langle\psi_t(s)| \right) \left(H(s) - i\hbar \tfrac{d}{ds} \right) |\psi_t(s)\rangle \right\| ds, \ \pi/2 \right\} \right).$$
(22)

The norm under the integral tracks the amount by which the trial state does not satisfy the Schrödinger equation. For the choice $H(t) = H_c(t) + H_e(t) + V(t)$, $|\psi(t)\rangle = |\psi_t(t)\rangle$, and $|\psi_t(t)\rangle = |\psi_n(t)\rangle$, the norm reduces to $\Delta(s)$, whence (19).

To explore the performance of (19), we have computed (22) numerically for a model system. The system is a particle in a step potential in one dimension, which splits an incoming wave packet into a transmitted and reflected wave packet, similar to a Hadamard gate. But here the potential models the interaction of the computer and environment. The initial state is a Gaussian wave packet. The trial state is a Gaussian wave packet with position and momentum equal to the classical motion across the potential step, neglects the reflected wave packet, but includes the spreading of the transmitted packet ("quasi-free particle"). Table II maps the computer dynamics onto the wave-packet dynamics. Figure 4 shows the result for incident kinetic energy 100 times the energy of the potential step, and initial wave-packet width 1/10 of the initial distance from the step. It is a sample from 15 computations, varying in initial energy and width. When the wave packet reaches the step and splits into a large-amplitude transmitted and small-amplitude reflected component, the overlap drops from 1.000 to 0.997; when the transmitted component is well separated from the reflected component, the overlap rises back to 1.000. In comparison, the


Figure 3. Lower bound τ for the decoherence time t_d (schematic). The upper curve is $F(t)$, assumed to obey $F(t) = e^{-t/t_d}$ for large t (definition of t_d); the lower curve is $F_-(t)$.

Table II. Wave-packet model (WP) of the dynamics of a quantum computer (QC).

	QC	WP
Isolated computer: • Hamiltonian • State	$H_c(t) + H_c(t)$ $\|\psi_n(t)\rangle$	Quasi-free particle $\|\psi_t(t)\rangle$
Nonisolated computer: • Hamiltonian • State	$H_c(t) + H_c(t) + V(t)$ $\|\psi_i(t)\rangle$	Particle in step potential $\|\psi(t)\rangle$
Fidelity: • Lower bound • Lower bound = 1 if • Lower bound ≈ 1 if	Eq. (19) $V(t) = 0$ Energy of $\|\psi_n(0)\rangle$ is large	Eq. (22) Potential step = 0 Traversal of step is fast

lower bound drops to 0.95 when the wave packet reaches the step, and remains at that value because the integral in (22) accumulates all departures of the trial state from the exact state: a temporary departure produces a permanent drop. The bound does not drop further because the trial state is asymptotically exact on the far side of the step. Thus the bound performs very well. It performs well whenever the wave packet traverses the step fast, i.e., the arrival time at the step is long and the residence time at the step is short.

Fast traversal of the step translates into a large energy of the computer state compared to the interaction energy. Thus the probability of a gate to fail at time t_g, $p(t_g)$, is small if the energy supplied to switch the gate, E_g, is large: $(\partial p / \partial E_g)_{t_g} < 0$. This agrees with the recent estimate, $p \sim \hbar/(t_g E_g)$, of the smallest achievable failure probability, or minimum energy required to switch the gate, under isolation [21]. Our result is a first step to extend the estimate to the case where the computer is not isolated.


Figure 4. Numerical evaluation of inequality (22). Left: overlap of the exact state and the trial state [left-hand side of (22)]. Right: lower bound of the overlap [right-hand side of (22)]. Time is in units of $\hbar/$(energy of the potential step).

CONCLUSIONS

We have reviewed quantum computing as an interplay of deterministic time evolution and probabilistic measurement outcomes, constructive and destructive interference, open and closed system dynamics, and theory and experiment. We have developed an inequality, Eq. (21), which decomposes the failure probability of the computer, due to imperfect isolation, into a part that maintains coherence, $|\psi_n(s)\rangle$, and the part that causes decoherence, $V(s)$. The inequality reduces the estimation of failure probabilities to a computationally easy problem for a wide range of computer designs. In an exploratory model computation, we find that the failure probability can be made small by making the energy stored in the computer large. This suggests there may be a variety of ways, to be uncovered yet, to minimize decoherence at given $V(s)$.

REFERENCES

1. B. Brezger et al., Phys. Rev. Lett. **88**, 100404 (2002).
2. P. Pfeifer, in *McGraw-Hill Yearbook of Science and Technology 2002* (McGraw-Hill, New York, 2001), pp. 294-298.
3. G.P. Berman, G.D. Doolen, R. Mainieri, and V.I. Tsifrinovich, *Introduction to Quantum Computers* (World Scientific, Singapore, 1998).
4. J. Gruska, *Quantum Computing* (McGraw-Hill, New York, 1999).
5. A.O. Pittenger, *An Introduction to Quantum Computing Algorithms* (Birkhäuser, Boston, 2000).
6. M.A. Nielsen and I.L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
7. C. Macchiavello, G.M. Palma, and A. Zeilinger (eds.), *Quantum Computation and Quantum Information Theory* (World Scientific, Singapore, 2001).
8. Special issue "Experimental Proposals for Quantum Computation," Fortschr. Phys. **48**, 769-1138 (2000).
9. Special issue "Quantum Computation for Physical Modeling," Comput. Phys. Comm. **146**, 277-344 (2002).
10. D. Gottesman and H.-K. Lo, Phys. Today **53** (11), 22 (2000).
11. P. Shor, in *Proceedings of the 35th Annual Symposium on the Foundations of Computer Science*, edited by S. Goldwasser (IEEE Computer Society, Los Alamitos, CA, 1994), pp. 124-134.
12. A. Ekert and R. Jozsa, Rev. Mod. Phys. **68**, 733 (1996).
13. S. Hallgren, in *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, edited by J. Reif (ACM Press, New York, 2002), pp. 653-658.
14 E. Knill, R. Laflamme, R. Martinez, and C.-H. Tseng, Nature **404**, 368 (2000).
15. L.M.K. Vandersypen et al., Nature **414**, 883 (2001).
16. C.A. Sackett et al., Nature **404**, 256 (2000).
17. A. Rauschenbeutel et al., Science **288**, 2024 (2000).
18. A.J. Leggett, Science **296**, 861 (2002); D. Vion et al., ibid., 886; Y. Yu et al., ibid., 889.
19. P. Pfeifer and J. Fröhlich, Rev. Mod. Phys. **67**, 759 (1995).
20. P. Pfeifer, Phys. Rev. Lett. **70**, 3365 (1993); **71**, 306 (1993).
21. J. Gea-Banacloche, Phys. Rev. Lett. **89**, 217901 (2002).